arXiv:1511.08635v1 [cs.PF] 27 Nov 2015

HPA: An Opportunistic Approach to Embedded Energy Efficiency

Baptiste Delporte and Roberto Rigamonti and Alberto Dassatti
Reconfigurable and Embedded Digital Systems Institute
REDS HEIG-VD, School of Business and Engineering Vaud
HES-SO, University of Applied Sciences Western Switzerland
Email: name.surname@heig-vd.ch

***Abstract*—Reducing energy consumption is a challenge that is faced on a daily basis by teams from the High-Performance Computing as well as the Embedded domain. This issue is mostly attacked from an hardware perspective, by devising architectures that put energy efficiency as a primary target, often at the cost of processing power. Lately, computing platforms have become more and more heterogeneous, but the exploitation of these additional capabilities is so complex from the application developer's perspective that they are left unused most of the time, resulting therefore in a supplemental waste of energy rather than in faster processing times.**

In this paper we present a transparent, on-the-fly optimization scheme that allows a generic application to automatically exploit the available computing units to partition its computational load. We have called our approach Heterogeneous Platform Accelerator (HPA). The idea is to use profiling to automatically select a computing-intensive candidate for acceleration, and then distribute the computations to the different units by off-loading blocks of code to them.

Using an NVIDIA Jetson TK1 board, we demonstrate that not only HPA results in faster processing speed, but also in a considerable reduction in the total energy absorbed.

I. Introduction

The energy consumption problem is common to the whole Computer Science world, as it is one of the major limiting factor to more powerful devices — especially embedded ones, which would otherwise run out of battery very quickly [1] — as well as one of the major sources of expenses (and pollution) for big data centers [2], [3]. Solutions to it range from accepting a performance reduction in return for longer battery life, as in the Intel Atom processor [4], to radical relocations of large data centers in cold regions [5].

In parallel to this drift towards greener forms of computing, the last few years saw a general trend in the direction of heterogeneous computing platforms, mostly issued from the acknowledgement of the limitations imposed by physics and technology to the pursuit of ever faster devices [6], [7], [8]. While having multiple, more energy-efficient accelerators on the same board — or even the same chip — should have led to increased performances and reduced power absorption (as each unit would be used in an optimal way, for instance, highly-parallel units for highly-parallel tasks), reality shows that the former target is only attained for very specific software customized for the accelerator, while the latter is rarely achieved due to the presence of a largely unused — but often leaking

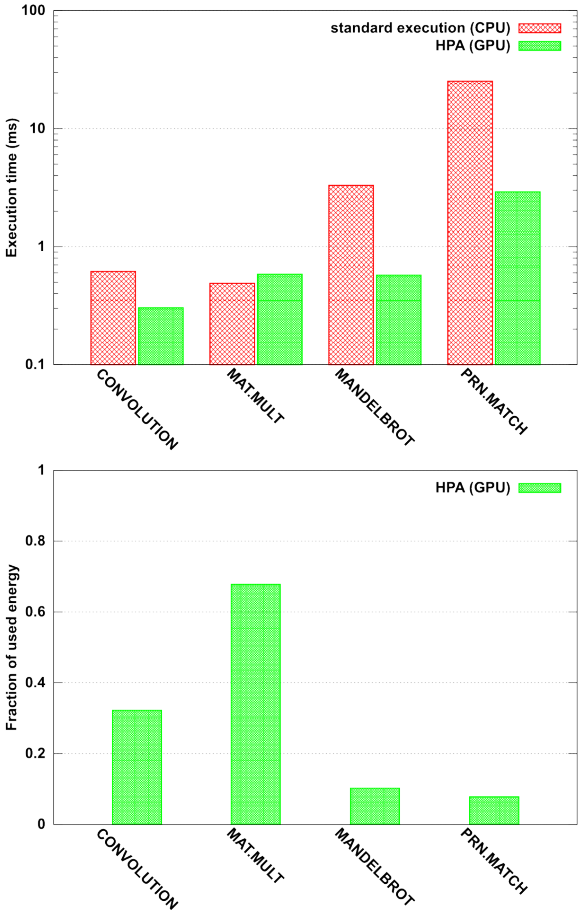

Fig. 1. **(top)** Execution time (in milliseconds over a logarithmic scale) of four test algorithms in the two situations we consider: standard CPU execution, and execution on the GPU in the context of our framework. We can see that HPA largely outperforms CPU execution in all but one test, MAT.MULT. However, if we consider the fraction of energy absorbed with respect to standard CPU execution **(bottom)**, we see that the execution in the context of HPA is considerably more efficient even when it takes more time to complete the computations.

power — resources. The problem this time lies on the software side: writing software for an inhomogeneous, ever changing set of targets is difficult, expensive, and requires skilled and motivated developers and constant maintenance. This fact contrasts with the very same nature of software, which usually evolves at a slow pace [9], and has heavily influenced current programming styles; As an example, Android developers are accustomed to not knowing on which platform the software they are writing is going to be executed. Indeed, most of the optimizations in JAVA, C#, Swift and other widespread programming languages are left at run-time, when the code is first execute and information about the target are available, generating in this way an optimized bytecode. However, this strategy reduces the scope of the optimizations available to the developer, forcing him to produce a generic code that is unlikely to fully exploit hardware capabilities. To make software capable of dealing with an heterogeneous environment and broader improvement opportunities, a reasonable approach seems to be automation [10], [11], [12].

In this paper we present an optimization system, called Heterogeneous Platform Accelerator (HPA), that automatically detects and delegates computing-intensive tasks to a dedicated accelerator, taking as an example of such accelerator the GPU chip available on an NVIDIA Jetson TK1 board. Our system grounds on the LLVM's Just-In-Time (JIT) compiler MCJIT [13], [14], [15]: executed code chunks are periodically analyzed by the perf_event [16] performance monitor and, once a particular function is classified as computing-intensive, it is ran through the Polly tool [17] to investigate whether it is parallelizable. If this is the case, the function is off-loaded to the GPU. The big advantage of this approach is that everything is totally transparent to the developer: he does not have to know on which platform his software will be deployed, as the potential accelerators will be discovered at run-time, nor he has to add any markings or adopt a particular library to allow the optimization. Moreover, the optimization is dynamic: should the performance analyzer detect that the GPU is slower than the CPU for a given task, or that it is desirable to leave the GPU to another function that would make better use of it, the system can revert its operations and recall back the currently off-loaded code. The unique limitation to the strategy we propose is the amount of intelligence and previous knowledge we are willing to put in the optimization scheme.

II. Related Work

In the last few years a remarkable increase in the interest about heterogeneous platforms has been observed; GPUs, by their nature and capabilities, have been deep in the center of this small revolution. Whilst a large number of proposals have focused, with good results, on devising a language layer that could ease the development on such platforms [18], [19], [20], relatively few implementations that automatically perform a multi-target optimization appeared.

An interesting approach, called BAAR, is presented in [10], [11]. The code to be executed is at first statically analyzed using Polly [17], a state-of-the-art polyhedral optimizer for automatic parallelization. If Polly detects that a function deserves being optimized, its code is compiled with the Intel's compiler and off-loaded to an Intel Xeon Phi board. All data transfers are dealt with using a software layer that handles them using MPI. While captivating results are given, the major drawback of this approach is that it lacks workload adaptation: the analysis is performed at application's startup, and it does not account for changes in the context of execution. The functions to be off-loaded are selected according to a metric that accounts for the number and type of operations to be performed, but this is only a rough indicator of the relative weight of the function in the context of the whole program's execution. As a consequence, this strategy is not reactive with respect to the user input, and this could lead to suboptimal choices — consider, for instance, the multiplication of a matrix whose size is user-specified.

Another proposal, named StarPU [21], provides an API and a pragma-based environment that, coupled with a run-time scheduler for heterogeneous hardware, composes a complete solution. While the main focus of the project is CPU/GPU systems, it could be extended to less standard systems. However, the developer is requested to learn a new API and foresee which parts of the code should be optimized. This analysis is dependent on both the input and the available resources and thus, despite the additional efforts required to the programmers, will not always lead to optimal results.

An approach which is not restricted to a single target type is SoSOC [22]: it consists of a library that presents a friendly interface to the programmer and allows functions to be dispatched to a set of targets based on either the developer's wishes or some statistics computed during early runs. This solution adds some dynamicity with respect to StarPU, but the developer not only has to learn (yet) another library, but also someone has to provide handcrafted code for any specialized unit of interest. This is a considerable waste of time and resources, and limits the applicability of the system to the restricted subset of architectures directly supported by the development team.

VPE [12] represents an evolution of SoSOC: its focus is on transparency, which means that the developer does not even have to be aware that his code is going to be accelerated. Starting by executing the code in an LLVM's JIT-based framework, VPE detects which functions are computing-intensive, and off-loads them to a remote unit. The code deployed is the very same code that is executed on the CPU, therefore no effort is requested to develop custom implementations. Results on a board based on the TI-DM3730 chip, which features a C64+ DSP processor, show gains up to $32\times$ in performance. Nonetheless, the implementation is heavily customized for the examined board, since LLVM does not provide a backend for this platform, and therefore a set of scripts has to compile ahead-of-time the functions using the TI proprietary compiler. When a candidate function is found, VPE off-loads its operations to the DSP by executing the previously-compiled code on the data that have been allocated in a shared memory region. In this paper we adopt a similar

approach, but with some relevant differences. Besides the different target family — a GPU — and the availability of a backend — that allows us to compile on-the-fly the portions of code of interest —, we put a strong bias towards energy efficiency, investigating how the optimization impacts on the power absorption. Moreover, we introduce an additional step in which computing-intensive functions are analyzed to detect the possibility of parallelizing them: indeed, it would be of little value to off-load sequential code to the 192-cores GPU of the Jetson TK1 board. Finally, we quantitatively investigate how the optimizations allowed by the use of a JIT-compiler in place of ahead-of-time compilation impact, in the context of our optimization scheme, both processing speed and power consumption.

III. Proposed Approach

The fundamental building block of our proposal is the LLVM's JIT compiler. LLVM is an alternative to the widely known GCC compiler with a neater separation between the front-end, the optimization, and the back-end steps [14]. This separation is largely due to the adoption of an enriched assembly language, called Intermediate Representation (IR), acting as a shared language between the different stages [13]. LLVM's community is very active, and recently a new JIT engine, called ORC, has been proposed. The previous engine, called MCJIT, presents indeed a serious issue: it has been designed to operate on modules — that is, aggregates of functions —, and once a module has been finalized (a mandatory step for execution), it is not possible to modify its code anymore. ORC solves this problem by allowing on-the-fly changes; however, it is currently under development and available for the x86_64 architecture only. For this reason, we have adopted the old JIT interface and implemented the transition across accelerators using a caller mechanism akin to that of [12]. In particular, we analyze at application startup which functions are available in the code — we automatically detect and exclude all system calls and I/O-based functions from our analysis, as we cannot optimize them with out approach —, and we replace function invocations with a caller that, when the function is not off-loaded, simply executes the desired function on the CPU via a function pointer. Once a function is selected for off-loading, we alter the function pointer to make it point to the function ready to be executed on the GPU. Please note that this caller overhead will be removed once ORC will be released for a broader selection of platforms.

Detecting whether a function deserves off-loading on the GPU or not is a two-step process: at first, functions that perform heavy computations are identified by using perf_event [16]. perf_event collects very detailed statistics about software and hardware counter, and permits us to easily identify performance bottlenecks. In the context of this paper we rely on the CPU usage alone to identify functions which might benefit from being off-loaded, but many more optimizations could be devised; For instance, inspecting memory access patterns could suggest an improved layout for a data structure that reduces cache misses [23]. After having obtained a sorted list of functions candidate for acceleration, we inspect each of them sequentially, checking with Polly [17] whether it is parallelizable or not. Polly operates at the IR level and starts by translating the code to optimize to a polyhedral representation [24]. It then detects the Static Controls Parts (SCoPs) [10] of the code matching a specific canonical form, and via a set of LLVM passes it performs the optimization and detects parallelism. In particular, it places OpenMP [25] markings around parallelizable blocks of code. We detect those markings and use them to guide our choices. For the sake of simplicity, we operate at function level, off-loading entire functions to the accelerator. Our approach is, however, independent with respect to the choice of the scale, so we could as well operate at the basic-block level [13] — which could be interesting, for instance, for a multi-threaded function.

Once we have selected a function to off-load to the GPU, we generate on-the-fly the PTX code [26] that is sent to the GPU by using the LLVM's backend. Data are transferred to an accessible memory region using the dedicated CUDA instructions, and the execution is started. The results are finally transferred back once the computation is finished. Although CUDA 6.5 offers an Unified Memory scheme [27], we experimentally verified that its performances are considerably worse than individual cudaMemcpy() operations. We have therefore opted for a manual transfer of the parameters, taking the transfer time into account in our measurements. Please note that, once the memory sharing will be a viable option, the performance of HPA will be further improved.

The architecture of HPA is depicted by Fig. 2.

IV. Results

We validate our proposal using an NVIDIA Jetson TK1 board, which features a 4+1 ARM Cortex-A15 32-bit processor with a 192-cores GK20A (Kepler) 852MHz GPU. This GPU, with respect to its predecessors, has a particular focus on energy efficiency. The installed Linux distribution is a Ubuntu 14.04 with a patched kernel distributed in the L4T (Linux for Tegra) package. We have set the CPU power governor to "performance" to guarantee the CPU the best possible performance. All the code we use has been compiled using Clang with strong optimization turned on (*-O3*).

We perform experiments with a test set similar to the one used in [12]: we considered a set of four algorithms inspired by the Computer Language Benchmarks Game[1], namely 2D convolution with a square kernel matrix (CONVOLUTION), multiplication of two square matrices (MAT.MULT), Mandelbrot set generation (MANDELBROT), and search of a nucleotidic pattern in an input DNA sequence (PRN.MATCH). Contrary to [12], given the nature of the adopted accelerator, we put no limitations on the use of floating point numbers. To provide a fair comparison, we fixed the amount of data to process and manually parallelized the CPU implementation of the algorithms to fully exploit the available cores.

[1] http://benchmarksgame.alioth.debian.org

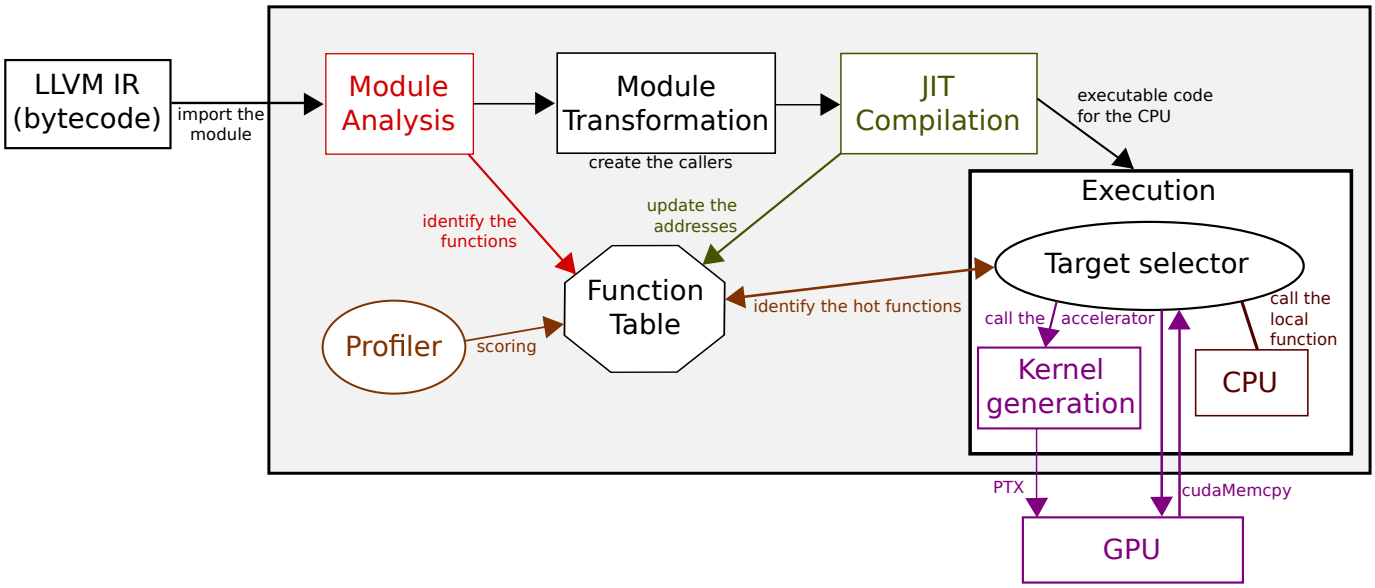

Fig. 2. Architecture of the HPA framework. An input program, in LLVM's IR bytecode format, is at first analyzed. Then, functions are identified and callers are created for all the potential candidates to off-loading. After the JIT-compilation step, the code is executed and profiling information are acquired. Once HPA detects that a given function is worth off-loading to the GPU, it transfers the required data and loads (or generates, if this is the first time the function is invoked) the PTX code that will be executed on the accelerator.

To estimate power consumption, we attached a shunt resistor to the ground power line and performed current measurements by using a Tektronix MSO5104B Mixed Signal Oscilloscope. We recorded the voltage over extensive periods and then computed the average and standard deviation values for each explored case.

Figure 1(top) shows the time required for the execution of each test in the two situations we consider, that is, standard CPU execution and execution on the GPU in the HPA framework. Executing on the GPU, even if the code is automatically generated and not hand-crafted for it, results in higher performance in all but the MAT.MULT case. This latter case could be due to a non-GPU-friendly implementation of the standard algorithm, and exposes the major weakness of our approach: as the code we execute on the accelerator was initially conceived for a standard CPU, we have little hope of achieving the performance of a carefully engineered CUDA algorithm. In our opinion, however, this drawback is largely offset by the complete transparency we offer the developer. If the system detects a lower-than-expected performance, it can at any time revert its choices and obtain as worst-case scenario the same performance the original code was supposed to achieve. Moreover, if we consider the power absorption depicted by Fig. 3, we see that, in our experiments, HPA is more energy efficient in all the cases, even when it takes longer to get to the final result as it is the case for MAT.MULT. Therefore, we could easily integrate an external policy that dictates which option to favor the most — best performance, best energy efficiency, or a combination of the two — and adapt the algorithm's behavior according to the run-time measurements. The combined effects of shorted processing times and lower power consumption are depicted by Fig. 1(bottom),

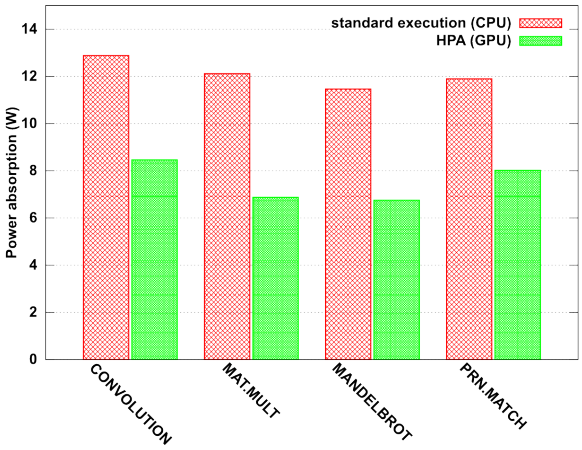

[t]

Fig. 3. Power absorption in the case of the four considered algorithms for both standard CPU execution and execution on the GPU in the context of our framework.

where the energy consumption of HPA is normalized over the consumption in the standard CPU case, while detailed results about execution time and power absorption are reported in Tab. I.

We have further investigated the MAT.MULT experiment to see the behavior of the system as a function of matrix size. Figure 4 depicts the execution time, the power consumption and the fraction of energy power used by the GPU with respect to standard CPU execution for a large set of matrix sizes. We can see that it is more convenient to operate on the CPU for matrices smaller than 200×200. Indeed, the data transfer we perform and the overhead introduced by the performance analyzer exceed by several orders of magnitude

the computation time, and therefore it can be up to $30\times$ more expensive, in energetic terms, to perform these multiplications on the GPU. For bigger sizes, despite the slightly higher execution time, the reduced power absorption makes the HPA alternative interesting.

We then developed an image-processing demo where different 2D filters are applied on the frames of a 1280×720 MPEG-4-encoded video. The image decompression and visualization tasks are performed using the Tegra version of the OpenCV library to avoid saturating the CPU with operations unrelated to the image-processing task itself, though this slightly penalizes HPA as it reduces the GPU resources available for the convolutions. In the comparison, we also considered the convolution operation performed by using the Tegra-OpenCV library: this version of the convolution is indeed optimized for the accelerator, and we expected it to outperform the HPA version. Moreover, we wanted to check whether using a JIT compiler — and thus being capable of adapting the code on-the-fly to the input data — was giving us some added value. We therefore JIT-compiled the module including the Constant Propagation optimization, and performed our measures using the three filters depicted by Fig. 6 (*emboss*, *Sobel* y-derivative, and *sharpen*) that exhibit different structures — in particular, different sparsity. The results are reported in Fig. 5 and detailed in Tab. II. Unsurprisingly, operating on the GPU resulted in significantly faster computations and thus, given a lower power absorption, a considerable reduction in the required energy with respect to CPU execution. The comparison with the Tegra-OpenCV implementation of the convolution operator is interesting: we can see that the frame rate of the hand-optimized convolution operator is lower than the one achieved by HPA for both the emboss and the Sobel filter. This could be due to the JIT compiler performing some small optimizations, and the genericity of the OpenCV version of the convolution preventing it from adopting architecture-dependent shortcuts. Also, we can see that, in the sharpening filter case, the OpenCV version can partially benefit of the presence of several zero entries in the filter matrix, but still it cannot outperform the HPA results when Constant Propagation is used in the JIT compilation. This latter, can indeed take advantage of the presence of elements with opposed sign, such as 1 and -1, and exploit this information to further speed-up computations. An example of these optimizations can be seen in Fig. 7, which has on the left side the generic convolution code that is normally executed by HPA and on the right side the code produced by JIT's Constant Propagation in the sharpening filter's case. These results support our claim that higher efficiency, in both computational and energetic terms, can be achieved by casting a problem in the HPA-JIT framework.

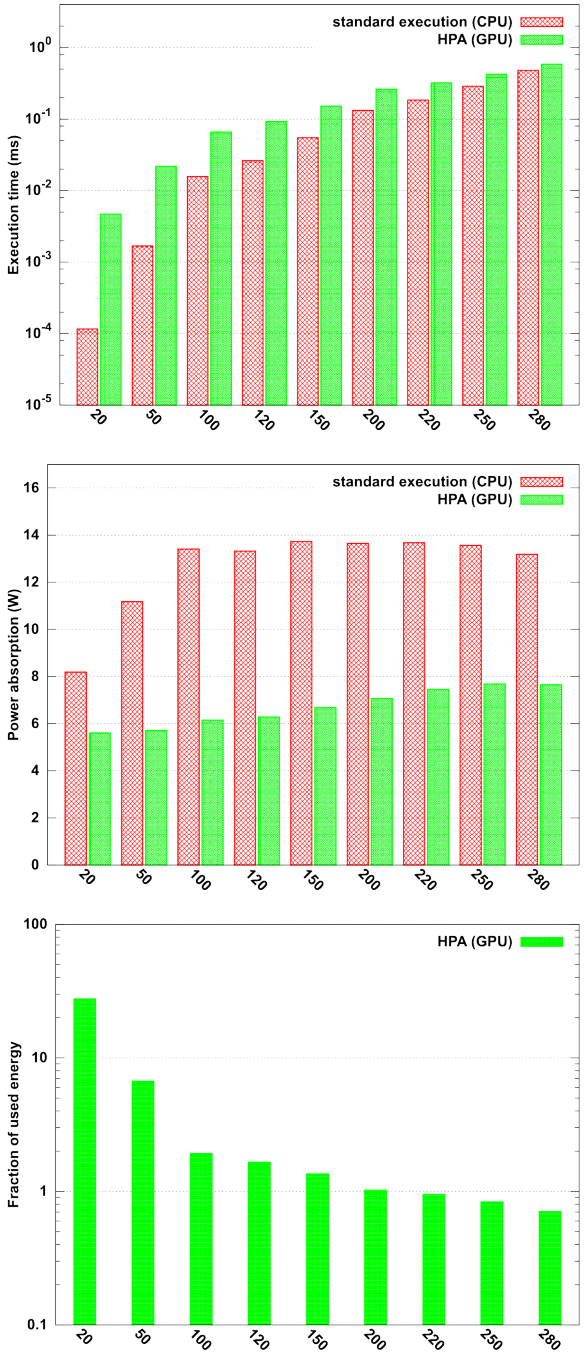

Fig. 4. **(top)** Execution time, **(center)** power absorption, and **(bottom)** fraction of energy with respect to the standard CPU execution for the MAT.MULT test as a function of matrix size. The graphs shows that it is orders of magnitude slower to execute the multiplication on the GPU when small matrices are involved; Indeed, both the data transfer time and the profiler's overhead exceed any potential gain that could derive by performing the multiplication on the GPU. Operating on the GPU becomes interesting, at least from an energetic stance, for matrices bigger than 200×200.

V. Conclusion

In this paper we presented an opportunistic strategy to increase energy efficiency by automatically off-loading computational intensive fragments of easily-parallelizable code to a GPU accelerator. As a side effect, we get a significant

TABLE I

EXECUTION TIME (IN MS) AND POWER ABSORPTION (IN W) FOR THE FOUR CONSIDERED ALGORITHMS. THE NUMBER REPORTED AFTER THE ± SIGN REPRESENTS ONE STANDARD DEVIATION. WITH "STANDARD EXECUTION (CPU)" WE INDICATE THE PLAIN EXECUTION OF THE ALGORITHM ON THE FOUR CORES OF THE ARM CPU WITH NO PERFORMANCE COLLECTION UNDERGOING, WHILE WITH "HPA (GPU)" WE INDICATE THE VERY SAME CODE BUT RUNNING ON THE GPU IN THE HPA FRAMEWORK

Algorithm	**standard execution (CPU)**		**HPA (GPU)**	
	Exec. time [ms]	Power cons. [W]	Exec. time [ms]	Power cons. [W]
CONVOLUTION	0.61 ± 0.001	12.88 ± 0.008	0.30 ± 0.043	8.46 ± 0.029
MAT.MULT	0.49 ± 0.015	12.11 ± 0.001	0.58 ± 0.088	6.87 ± 0.001
MANDELBROT	3.30 ± 0.018	11.46 ± 0.013	0.57 ± 0.131	6.75 ± 0.016
PRN.MATCH	25.20 ± 0.008	11.89 ± 0.005	2.91 ± 0.280	8.01 ± 0.006

TABLE II

FRAME RATE (IN FPS) AND POWER ABSORPTION (IN W) FOR THE IMAGE PROCESSING TEST. THE NUMBER REPORTED AFTER THE ± SIGN REPRESENTS ONE STANDARD DEVIATION. WITH "STANDARD EXECUTION (CPU)" WE INDICATE THE PLAIN EXECUTION OF THE ALGORITHM ON THE FOUR CORES OF THE ARM CPU WITH NO PERFORMANCE COLLECTION UNDERGOING, WITH "HPA (GPU)" WE INDICATE THE VERY SAME CODE BUT RUNNING ON THE GPU IN THE HPA FRAMEWORK, WITH "TEGRA-OPENCV (GPU)" WE INDICATE THE EXECUTION ON THE GPU OF THE TEGRA-OPENCV OPTIMIZED CONVOLUTION CODE, AND WITH "HPA + CONSTANT PROPAGATION (GPU)" THE HPA VERSION BUT WITH THE CONSTANT PROPAGATION OPTIMIZATION OF THE JIT FRAMEWORK ACTIVE

Execution target	**Frame rate [fps]**			**Power absorption [W]**		
	Emboss	Sobel	Sharpen	Emboss	Sobel	Sharpen
standard execution (CPU)	23.97 ± 0.61	23.68 ± 0.61	23.07 ± 0.78	9.00 ± 0.015	11.02 ± 0.027	11.35 ± 0.094
HPA (GPU)	53.84 ± 6.80	54.25 ± 7.69	54.47 ± 7.81	8.04 ± 0.022	9.62 ± 0.018	9.69 ± 0.043
Tegra-OpenCV (GPU)	48.88 ± 3.12	50.06 ± 2.91	65.48 ± 7.15	9.39 ± 0.213	9.38 ± 0.093	9.55 ± 0.009
HPA + Constant Propagation (GPU)	63.92 ± 8.45	64.02 ± 9.03	70.02 ± 7.90	8.22 ± 0.007	7.99 ± 0.023	9.58 ± 0.104

increment in the overall performances, as not only these computations are performed faster, but also the main CPU load is relieved and thus it can accomplish further tasks while waiting for the computations to terminate. We supported our claims with thorough experiments on several algorithms and an image processing task.

As future work we will focus on defining other optimization strategies and investigate platforms with a higher degree of heterogeneity, choosing at run-time the target that is expected to give the highest energy efficiency or fit best to a set of user-defined policies.

REFERENCES

[1] M.T. Schmitz and B.M. Al-Hashimi and P. Eles, *System-Level Design Techniques for Energy-Efficient Embedded Systems*. Kluwer Academic Publishers, 2004.

[2] J.G. Koomey, "Worldwide electricity used in data centers," *Environmental Research Letters*, 2008.

[3] S. Ruth, "Green IT - More Than a Three Percent Solution?," *IEEE Internet Computing*, 2009.

[4] B. Beavers, "The Story behind the Intel Atom Processor Success," *IEEE Design Test of Computers*, 2009.

[5] S. Hancock, "Iceland looks to serve the world," *BBC News*, 2009.

[6] S.A. Khan, *Digital Design of Signal Processing Systems: A Practical Approach*. Wiley, 2011.

[7] E.S. Chung and P.A. Milder and J.C. Hoe and K. Mai, "Single-Chip Heterogeneous Computing: Does the Future Include Custom Logic, FPGAs, and GPGPUs?," in *Proc. of the IEEE/ACM Int. Symposium on Microarchitecture*, 2010.

[8] R. Kumar and D.M. Tullsen and N.P. Jouppi and P. Ranganathan, "Heterogeneous Chip Multiprocessors," *Computer*, 2005.

[9] F.P. Brooks, *The Mythical Man-month (Anniversary Ed.)*. 1995.

[10] M. Damschen and C. Plessl, "Easy-to-use on-the-fly binary program acceleration on many-cores," in *Proc. Int. ASCS Workshop*, 2015.

[11] M. Damschen and H. Riebler and G. Vaz and C. Plessl, "Transparent Offloading of Computational Hotspots from Binary Code to Xeon Phi," in *Proc. of the Design, Automation & Test in Europe Conference & Exhibition*, 2015.

[12] B. Delporte and R. Rigamonti and A. Dassatti, "Toward Transparent Heterogeneous Systems," in *Submitted to the MULTIPROG-2016 Workshop*, 2015.

[13] C. Lattner and V. Adve, "LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation," in *Proc. of the CGO Symposium*, 2004.

[14] C. Lattner, "Introduction to the LLVM Compiler System," in *Proc. of the ACAT Workshop*, 2008.

[15] C. Lattner, "LLVM and Clang: Advancing Compiler Technology," in *Proc. of the FOSDEM*, 2011.

[16] V.M. Weaver, "Linux perf_event Features and Overhead," in *Proc. of the FastPath Workshop*, 2013.

[17] T. Grosser, A. Groesslinger, C. Lengauer, "Polly - Performing polyhedral optimizations on a low-level intermediate representation," *Parallel Processing Letters*, 2012.

[18] A. Danalis and G. Marin and C. McCurdy and J.S. Meredith and P.C. Roth and K. Spafford and V. Tipparaju and J.S. Vetter, "The Scalable Heterogeneous Computing (SHOC) Benchmark Suite," in *Proc. of the General-Purpose Computation on Graphics Processing Units Workshop*, 2010.

[19] J.E. Stone and D. Gohara and G. Shi, "OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems," *Computing in Science & Engineering*, 2010.

[20] G. Kyriazis, "Heterogeneous system architecture: A technical review," tech. rep., AMD, 2013.

[21] C. Augonnet and S. Thibault and R. Namyst and P.A. Wacrenier, "StarPU: A Unified Platform for Task Scheduling on Heterogeneous Multicore Architectures," *Concurrency and Computation: Practice & Experience*, 2011.

[22] O. Nasrallah and W. Luithardt and D. Rossier and A. Dassatti and J. Stadelmann and X. Blanc and N. Pazos and F. Sauser and S. Monnerat, "SOSoC, a Linux framework for System Optimization using System on Chip," in *Proc. of the IEEE System-on-Chip Conference*, 2013.

[23] T.M. Chilimbi and M.D. Hill and J.R. Larus, "Cache-Conscious Structure Layout," in *Proc. of the PLDI Conf.*, 1999.

[24] U. Bondhugula and A. Hartono and J. Ramanujam and P. Sadayappan, "A Practical Automatic Polyhedral Parallelizer and Locality Optimizer," *SIGPLAN Not.*, 2008.

[25] L. Dagum and R. Menon, "OpenMP: an industry standard API for shared-memory programming," *IEEE Computational Science Engineering*, 1998.

[26] NVIDIA Corporation, *PTX: Parallel Thread Execution*. 2008.

```
// Generic convolution
// Generated by LLVM NVPTX Back-End, using HPA
< ... >
// BB#1:                                // %if.end
    ld.param.u32  %r5, [convolution_param_2];
    ld.param.u32  %r4, [convolution_param_1];
    ld.param.u32  %r3, [convolution_param_0];
    add.s32       %r13, %r4, %r1;
    ld.u8         %r14, [%r13];
    ld.u32        %r15, [%r3];
    mul.lo.s32    %r16, %r14, %r15;
    ld.u8         %r17, [%r13+1];
    // ** The following block is optimized in the **
    // ** HPA + Constant Propagation version      **
    ld.u32        %r18, [%r3+4];
    ld.u8         %r19, [%r13+2];
    ld.u32        %r20, [%r3+8];
    add.s32       %r21, %r13, %r2;
    ld.u8         %r22, [%r21];
    ld.u32        %r23, [%r3+12];
    ld.u8         %r24, [%r21+1];
    ld.u32        %r25, [%r3+16];
    ld.u8         %r26, [%r21+2];
    ld.u32        %r27, [%r3+20];
    add.s32       %r28, %r21, %r2;
    ld.u8         %r29, [%r28];
    ld.u32        %r30, [%r3+24];
    ld.u8         %r31, [%r28+1];
    ld.u32        %r32, [%r3+28];
    ld.u8         %r33, [%r28+2];
    ld.u32        %r34, [%r3+32];
    mad.lo.s32    %r35, %r17, %r18, %r16;
    mad.lo.s32    %r36, %r19, %r20, %r35;
    mad.lo.s32    %r37, %r22, %r23, %r36;
    mad.lo.s32    %r38, %r24, %r25, %r37;
    mad.lo.s32    %r39, %r26, %r27, %r38;
    mad.lo.s32    %r40, %r29, %r30, %r39;
    mad.lo.s32    %r41, %r31, %r32, %r40;
    mad.lo.s32    %r42, %r33, %r34, %r41;
    // ** End of the replaced block **
    // 15 load operations
    // 2 additions
    // 8 "multiply and add" operations
< ... >
```

```
// Optimized SHARPEN convolution
// Generated by LLVM NVPTX Back-End, using HPA
< ... >
// BB#1:                                // %if.end
    ld.param.u32  %r4, [convolution_param_2];
    ld.param.u32  %r3, [convolution_param_1];
    add.s32       %r12, %r1, %r2;
    add.s32       %r13, %r3, %r12;
    ld.u8         %r14, [%r13];
    // ** The following block has been optimized **
    ld.u8         %r15, [%r13+1];
    // The compiler has detected the -1 in the matrix
    sub.s32       %r16, %r15, %r14;
    // ** End of the optimized block **
    // 1 load operation
    // 1 subtraction
< ... >
```

Fig. 7. PTX code generated by the LLVM back-end for the generic convolution code **(left)** and the version optimized by the JIT Constant Propagation optimization step **(right)**. The number of executed instruction is greatly reduced due to the particular shape of the adopted image kernel.

[27] M. Harris, "Unified memory in CUDA 6," in *GTC On-Demand*, 2013.

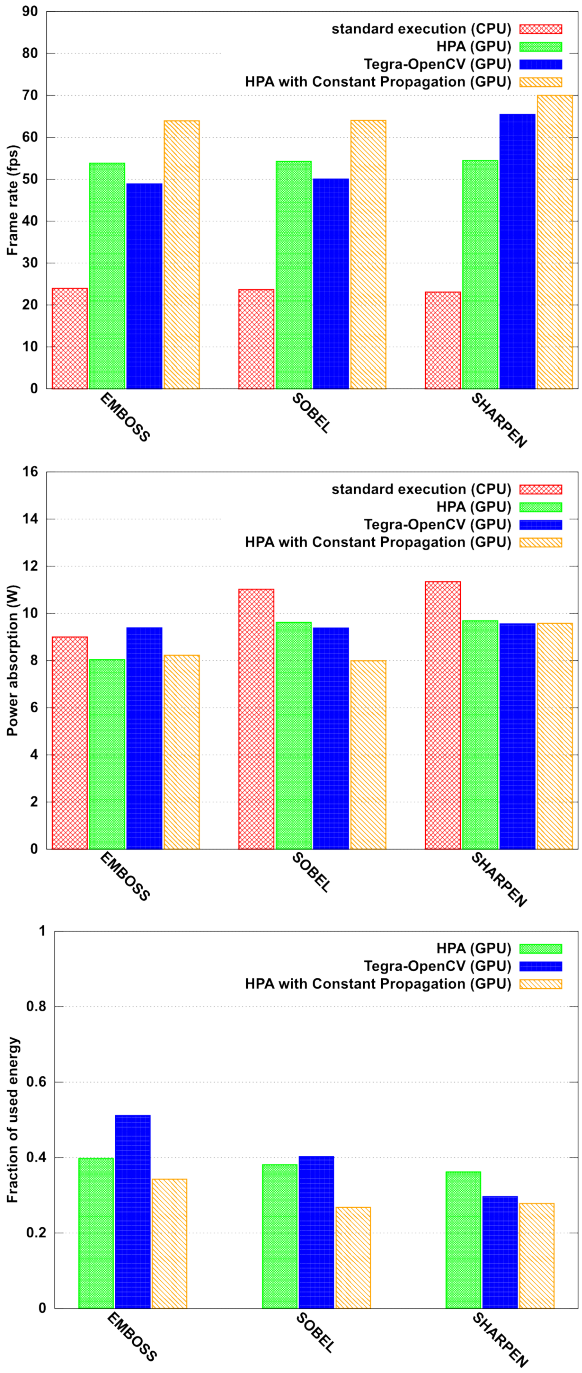

Fig. 5. **(top)** Execution time, **(center)** power absorption, and **(bottom)** fraction of energy with respect to the standard CPU execution for the image processing demo. Three different filter types are considered, each with a different sparsity degree. We can see that the Tegra-OpenCV implementation can partially exploit the sparsity, but not as much as the JIT-optimized version — which is unsurprising, as the OpenCV version still has to perform all the multiplications, while the JIT version drops all unnecessary operations.

`emboss`
$$\begin{bmatrix} -2 & -1 & 0 \\ -1 & 1 & 1 \\ 0 & 1 & 2 \end{bmatrix}$$

`Sobel`
$$\begin{bmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{bmatrix}$$

`sharpen`
$$\begin{bmatrix} 0 & 0 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Fig. 6. Structure of the three filters used in the image processing demo. We have chosen them as they present a different degree of sparsity, and we wanted to investigate the capabilities of a JIT-based framework to exploit this information when optimizing the computations.